REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

AD-A255 432

1. AGENCY USE ONLY

3. REPORT TYPE AND DATES COVERED
ANNUAL 15 Apr 91 TO 14 Apr 92

4. TITLE AND SUBTITLE
PSYCHOPHYSICAL ANALYSES OF PERCEPTUAL REPRESENTATIONS

5. FUNDING NUMBERS
AFOSR-90-0274
PE - 61103D
PR - 3484
TA - HS

6. AUTHOR(S)
Dr Irving Biederman

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Minnesota
Department of Psychology
75 East River Road
Minneapolis, MN 55455

8. PERFORMING ORGANIZATION REPORT NUMBER
4

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Dr John F. Tangney
AFOSR/NL
Building 410
Bolling AFB DC 20332-6448

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

DTIC
ELECTE
SEP 14 1992
S A D

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release;
distribution unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

Research during the year has been divided between studies at USC (Biederman and students) and Minnesota. Our research continues to focus on linking early sensory representations to higher-level perceptual representations. Studies outlined below have examined the sensory/perceptual "middle ground" in object recognition, depth perception, reading, and auditory perception. Several of our studies have used ideal-observer analysis. The ideal-observer approach provides a means for quantifying the information available to perception and for evaluating the effectiveness with which humans use that information.

92 9 11 02

92-25079

12 Pg

14. SUBJECT TERMS

15. NUMBER OF PAGES

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
(U)	(U)	(U)	(UL)

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)

ANNUAL TECHNICAL REPORT

Submitted to: Air Force Office of Scientific Research
AFOSR/NL
Building 410
Bolling A.F.B., D.C. 20332-6448

Dr. John Tangney, Program Manager

Title: *PSYCHOPHYSICAL ANALYSES OF PERCEPTUAL REPRESENTATIONS*

Grant No. AFOSR-90-0274

Period Covered: April 16, 1991 - April 15, 1992

Submitting Organization: University of Minnesota
Office of Research & Technology Transfer Admin.
1919 University Ave.
St. Paul, MN 55104

PI:[1] Professor Gordon E. Legge
(612) 625-0846

Gordon E. Legge
023-42-1595

Date of Report: 18 August 1992

[1]An amendment is pending whereby funds from this grant will be subcontracted to the University of California for research under the direction of Professor Irving Biederman.

PROGRESS REPORT: 1991-1992

Introduction

Research during the year has been divided between studies at USC (Biederman and students) and Minnesota. Our research continues to focus on linking early sensory representations to higher-level perceptual representations. For this reason, we refer to our Center informally as the "Middle Kingdom." Studies outlined below have examined the sensory/perceptual "middle ground" in object recognition, depth perception, reading, and auditory perception. Several of our studies have used ideal-observer analysis. The ideal-observer approach provides a means for quantifying the information available to perception and for evaluating the effectiveness with which humans use that information.

In the following paragraphs, we describe many projects supported by the grant. A list of publications and conference presentations follows these descriptions.

Recognizing Depth-Rotated Objects: Evidence for 3D Viewpoint Invariance (Biederman & Gerhardstein, 1992)

Several recent reports have documented extraordinary difficulty in the recognition of images of certain kinds of unfamiliar 3D objects from a novel orientation in depth. The difficulty at specific orientations can be greatly reduced with practice at those orientations. If generally true, such a result would support the contention that the capacity to recognize everyday objects is a consequence of familiarity over a variety of viewpoints, in which separate visual representations (templates) are created for each experienced viewpoint. Such a theory would stand in contrast to invariant-parts theories of basic level object recognition which assume that a viewpoint invariant structural description (up to parts occlusion and accretion) can be created from a single view of many objects, whatever their familiarity. Three experiments are reported. The first revealed complete viewpoint invariance in the visual (not just name or concept) priming of novel images of familiar objects in that changes of up to 135 deg in depth resulted in virtually no reduction in the magnitude of facilitation of naming RTs. The second experiment showed that priming could be reduced if there was a change in the part descriptions from priming to primed trials. The third experiment employed unfamiliar objects composed of novel arrangements of volumes. Same-different judgments of sequentially presented images showed little cost of rotation in depth as long as the same invariant parts description could be activated. Together these results suggest that depth invariance can be readily achieved if the different stimuli activate distinctively different and viewpoint invariant (e.g., geon) representations. These two specifications may constitute the

defining perceptual conditions for the formation of basic (or entry) level categories.

Priming Objects with Single Volumes: Searching for the Representational Locus of Perceptual Priming. (Cooper and Biederman, 1992)

Subjects' latencies to name an object picture decrease with repeated presentations of the object (Bartram 1974). Two experiments were conducted to determine the representational level at which the perceptual portion of this priming occurs. Subjects named objects that could be preceded by a single volume prime (which could either be present or absent in the object) or a neutral line. No effect of prime type was found on object naming RTs or errors even when the objects' identities were made salient by displaying them beforehand. These results in combination with previous experimental data (Biederman & Cooper 1991) support a representational level specifying an object's convex components and their relations to one another as the locus of visual priming.

High Level Object Recognition Without an Inferior Temporal Lobe. (Biederman, Gerhardstein, Cooper and Nelson, 1992)

Seven individuals with unilateral temporal lobectomies (four left and three right), in which the anterior and medial regions of the inferior temporal lobe were removed, and 8 controls, performed two types of shape recognition tasks with briefly presented, lateralized line drawings of 3D objects. In a same-different task, the subjects judged whether line drawings of two objects, presented sequentially with an intervening mask, were the same or different in shape, disregarding differences in orientation up to 60 deg in depth. The objects were either familiar or nonsense objects. With the familiar objects, different trials were of different shaped exemplars with the same name. In the other task, subjects named familiar objects. In either task, the disadvantage of presenting an image to the lobectomized hemisphere, either initially or in a second priming block, was negligible. These results indicate that efficient high-level object recognition does not require the anterior and medial regions of the temporal lobe in the hemisphere that initially receives an image. Object recognition is either accomplished more posteriorly, perhaps at the temporal-occipital boundary, or by the remaining temporal lobe, through a completely efficient callosal connection.

To What Extent Can Matching Algorithms Based on Direct Outputs of Spatial Filters Account for Human Shape Recognition? (Fiser, Biederman, and Cooper)

After an initial filtering of the image, models of basic level

object recognition typically posit several intervening stages which create an intermediate representation that, in turn, activates an object class description. This research evaluated a highly successful face recognition system based on von der Malsburgs's Dynamic Link Architecture (DLA) theory as a model of human object recognition. The system consists of only two layers: The output of an array of Gabor kernels on multiple scales and orientations is mapped directly onto stored representations to achieve recognition, preserving the topographic relations among the outputs of Gabor kernels in the matching phase. The system attempted to recognize contour deleted versions or mirror reflections of a gallery of line drawings of common objects. System accuracy was quite high overall, however the performance of the system was qualitatively different from that evidenced by humans in real-time shape recognition tasks. Thus, although the system's filter output description may be appropriate for initial representation of information in the visual scene, it does not provide a good model of human object recognition. Modeling of human object recognition might require a structural description of shape that explicitly specifies the information that allows classification. The DLA system likely derives its recognition power from its capacity to represent precisely metric spatial relations for grey scale variation--something that people may not be able to use.

Object Recognition and Classification for Human and Ideal Observers. (Liu, Kersten & Knill, 1992)

We developed a novel paradigm for experimental studies of human object recognition (Liu, Kersten, and Knill, 1992; Liu, Knill, and Kersten, 1992). By computing the statistical efficiency of human observers relative to an ideal observer for an object classification task, we obtain an absolute measure of the ability of subjects to use the stimulus information for the task. The measured efficiencies enable us to make strong inferences about the architecture of the recognition systems used by human observers. In this paper, we measure the statistical efficiency with which human observers make simple classification judgments of randomly shaped thick wire objects. After training to 11 different views of an object, subjects were asked which of a pair of noisy views of the object best matched the learned object. Human statistical efficiencies relative to a 2D ideal which based its judgment on a simple template matching strategy exceeded 100%. These high efficiencies exclude models which are suboptimal relative to the 2D ideal, such as view based Hyper Basis Function interpolation models which only include the 2D spatial coordinates of object features in the input representation. Instead, the results indicate that 3D constraints, above and beyond those implicit in the 2D ideal, are incorporated in the recognition process. Moreover, object regularity (e.g. planarity, symmetry) dramatically improved the efficiency with which novel views of an object could be classified.

The Geometry of Shadows. (Knill, Mamasian & Kersten, 1992)

Shadows provide a strong source of information about the shapes of surfaces. Drawing on previous work on smooth occluding contours, we analyze the local geometric structure of shadow contours on smooth surfaces. We also consider the behavior of shadow contours on piece-wise smooth surfaces. Particular attention is paid to intrinsic shadows on a surface; that is, shadows created on a surface by the surface's own shape and placement relative to a light source. We analyze the invariants relating surface shape to the shapes and singularities of bounding contours of such shadow contours, including the singularities in the evolution of shadows on a surface as it is moved relative to a light source. We show that the results obtained for point sources of light generalize in a straightforward way to extended light sources, under the assumption that light sources are convex.

Spatial Layout from Cast Shadows. (Mamasian, Kersten & Knill, 1992)

When an object casts its shadow on a background surface, the distance separating the object from the shadow, as it appears in the image, provides information about the position of the object relative to the background. In comparison to other pictorial depth cues such as occlusion, however, shadows have received little attention. In this study, we investigate the perceived 3D motion of an object in the presence of a shadow. We provide a method to measure the influence of cast shadows on spatial layout perception, natural image constraints on shadow formation, and the interaction of cast shadows with other depth cues. The stimulus consists of a ball and the shadow it casts on the bottom of a surrounding box. The box is rendered in perspective projection. The ball is given an oscillating motion along a linear path of a fixed orientation in the image. The trajectory of the shadow is also linear, but its orientation is an independent variable, so that the maximum distance between ball and shadow varies from trial to trial. The perceived height of the ball relative to the bottom plane is assessed with the help of a displaceable landmark on one lateral side of the box. The results show that the greater the distance between ball and shadow, the higher the ball is perceived. The method described allows us to ask: What are the necessary characteristics of a "patch" in the image to provide spatial information compatible with a cast shadow? Such potential characteristics include the two following brightness constraints: the shadow region should be darker than its surround, and the contrast polarity should be conserved along the shadow boundary. In contrast with previous findings on shape from shadow, it seems that these brightness constraints can be violated. In particular, the coherency between object and shadow motions appears to be a stronger constraint than the shadow darkness or the contrast

polarity along the boundary.

Structure-from-Motion Based on Information at Surface Boundaries. (Thompson, Kersten & Knecht, 1992)

Existing computational models of structure-from-motion -- the appearance of three-dimensional motion generated by moving two-dimensional patterns -- are all based on variations of optical flow or feature point correspondences within the interior of single objects. Three separate phenomena provide strong evidence that in human vision, structure-from-motion is significantly affected by surface boundary cues. In the first, a rotating cylinder is seen, though no variation in optical flow exists across the apparent cylinder. In the second, the shape of the bounding contour of a moving pattern dominates the actual differential motion within the pattern. In the third, the appearance of independently moving objects changes significantly when the boundary between them becomes indistinct. We describe a simple computational model sufficient to account for these effects. The model is based on qualitative constraints relating possible object motions to patterns of flow, together with an understanding of the patterns of flow that can be discriminated in practice.

A Multi-Layer Approach to Segmentation and Interpolation. (Madarasmi, Kersten & Pong, 1992)

Computational methods for surface interpolation and segmentation often use smoothness processes to constrain the surface interpretation within statistically correlated regions and line processes to describe the discontinuities between these smooth regions. This edge-based approach to interpolation does not explicitly segment the data into meaningful regions and does not work well for segmenting images containing transparent regions. We present a multi-layer approach to the segmentation and interpolation problem which partitions the input image into separate layers, each corresponding to a smooth region in the image, and simultaneously fills in the missing data within each layer. The proposed multi-layer system can successfully segment both opaque and transparent images within a single computational framework. Thus, given one image with both opaque regions and transparent regions, the system will compute the appropriate segmentation without treating the two regions differently. The system is shown to be particularly appropriate for the stereo matching paradigm. Stereo matching, interpolation, and segmentation are performed simultaneously to achieve the correct correspondence for both opaque and transparent surfaces, bringing together the classical stereo correspondence theories for the two types of surfaces. The results from computer simulations for segmenting intensity images and for computing disparity in random-dot

stereograms and in real stereograms are presented.

The Perception of Surface Marking Contours and Surface Shape. (Knill, 1992)

We have continued our study of the role played by surface marking contours in the perception of surface shape. Surface marking contours are contours projected from extended markings (e.g. reflectance edges) on surfaces. Numerous demonstrations have shown the effectiveness of such contours in eliciting a perception of curved surface shape. The underlying hypothesis we have developed and are testing is that the visual system incorporates an assumption that surface markings are geodesic in the inference of surface shape from surface marking contours. A paper summarizing the mathematical analysis of what we have called a "geodesic constraint" as well as some psychophysical support for its psychological validity will appear in the Journal of the Optical Society of America. Our current efforts in this project are focused on developing computer simulations of a model which incorporates the geodesic constraint in the interpretation of surface shape from surface marking contours.

The Statistical Structure of Contours. (Knill)

We have begun a theoretical analysis of the statistical structure of contours in natural images. This has resulted in a number of surprising results concerning the probability distribution of corner angles in images and the assumptions required to explain such phenomenon as the perception of skew symmetries as oriented real symmetries. Furthermore we have shown that several previous models of planar surface orientation estimation from contour shape implicitly assume a fractal structure of contours. This has led us to formulate an improved model of the assumptions underlying the estimation of planar surface orientation from contour shape. Moreover, the observation that contours in natural images may have a fractal structure has implications for the coding of contours in the early visual system. We have begun an investigation into these issues.

Contour Shape Perception. (Knill)

We have begun to study the efficiency with which humans can discriminate various aspects of contour shape (curvature, corner angle, skewness, etc.). Besides answering questions about how the visual system codes contour shape, the results of this study will provide limits on the reliability with which the visual system can infer surface shape from image contours, whether they be occluding contours, shadow contours or surface marking contours. The study relies on a novel use of the ideal observer approach as it has been

applied to simple signal discrimination tasks. We have completed the mathematical analyses prerequisite to the application of the approach to psychophysical studies and are beginning to run experiments.

The Perception of Illuminant Direction and Shape from Shading. (Knill)

Most models of shape from shading require that the position of the dominant light source illuminating a surface be known. Several sources of information for light source direction, including the global statistical structure of surface shading and the shape of shadow contours, have been identified. We have begun a psychophysical investigation into questions about what information determines human perception of light source direction. In contrast to previous studies, we have developed an experimental paradigm which allows us to measure subjects' perception of light source direction indirectly through estimates of shape characteristics of surfaces. This allows us to tap into the perceptual estimation of light source direction actually involved in the estimation of shape from shading. Through pilot studies, we have perfected the technique and are beginning to run full-scale studies of perceptual light source estimation.

The Role of Color in Object Recognition. (Wurm, Legge, Isenberg & Luebker, 1992)

Does color improve object recognition? If so, is the improvement greater for blurred images where there is less shape information? Do people with low visual acuity benefit more from color than people with normal acuity? We addressed these questions in three experiments by comparing naming reaction times (RTs) for food objects displayed in four ways: achromatic or color, and blurred or unblurred. Normally sighted subjects had faster reaction-times with color that did not change significantly with blur. Low-vision subjects were also faster with color and the difference did not depend significantly on acuity. In two additional experiments, we asked if the faster RTs for color stimuli were related to objects' prototypicality or color diagnosticity. We conclude that color does improve object recognition and the mechanism is probably sensory rather than cognitive in origin.

Statistical Efficiency for Categorization of Curvature: Effects of Viewpoint Invariance. (Mansfield, Biederman, Knill & Legge, 1991)

We have measured the statisitical efficiency with which observers can classify curved contours (circular arcs of fixed arc length) into pretrained curvature categories. A categorization task was used for these measures purposefully in an attempt to reveal the

manner in which curvature is internally represented by the visual system. Subjects could perform the task with efficiencies as high as 80% if one of the curvature categories included zero curvature. When the categories were moved away from zero curvature efficiencies decreased by as much as 30%. As we reported previously, these data are consistent with the exploitation of viewpoint invariance differences in object recognition. We have also shown that the efficiencies measured in the categorization task cannot be accounted for by differences in the discriminability of curvature.

The Perceived Location of Binocular Depth Targets. (Mansfield, Akutsu & Legge, 1992)

Many schemes have been proposed that can account for the encoding of depth using binocular vision. In this project we have been considering how a binocular visual system might encode both depth and horizontal location. Such a process is necessary for performing alignment tasks in 3D, or for the veridical perception of object shape. In principle, for a known fixation distance, the information in the two eyes is sufficient to enable correct localization of objects in 3D. The depth of a feature is proportional to the difference in the 'local signs' in the views of each eye, whereas direction is related to the average of the 'local sign' values in each eye. Using a vernier alignment procedure, however, we have shown that if the images presented to each eye have different luminance contrasts, then the perceived direction of the stereo target is biased towards the view seen by the eye with higher contrast. This shift in perceived direction is consistent with the visual system choosing the "most-likely" visual direction given the information in each eye.

The Role of Font Information in Reading. (Klitz, Mansfield & Legge, 1992).

What is the role of font information in reading? Casual inspection of the printed material we read everyday shows that there is a variety of different fonts, some of which are very different from one another, and others that seem almost identical. Typically, reading is effortless regardless of the style of type used, and apart from odd occasions, we are quite unaware of font as we read. However, words written in a bold or itlaic font sometimes 'pop out' from a page of text, which might suggest that fonts are involved in a global analysis of the text page. We have used a reaction time task to examine the perceptual distinctiveness of pairs of fonts. Subjects were required to detect a 'target' region of text rendered in a different font from the 'background' text. The size (number of glyphs) in the target could be varied. The results show that for some pairs of fonts, the reaction time for detecting the target was

the same irrespective of the target size. These font pairs could be said to 'pop-out' from one another, implying that they were detected in a rapid global analysis of the page. For other pairs of fonts, the reaction times were dependent on the target size. Some pairs of fonts always produced long reaction times, and had large error rates. In these latter cases the font pairs were perceptually indistinct. These data allow us to speculate that the font information my be used in a fast global page analysis. Also, the reaction time data can be used as a measure of font distinctiveness which will provide a useful set of benchmarks for later experiments in this study.

Mr. Chips: An Ideal Observer Model of Reading. (Legge, 1992)

Existing models of reading do not explicitly specify how visual data are combined with other sources of information, nor do they explain how visual disorders affect reading. Ideal-observer models have been useful in vision because they are explicit in identifying sources of information and task constraints. The perceptual component of reading can be formalized as the interpretation of a string of stimulus symbols (text), sampled through a window whose position is determined by a sequence of saccades. An ideal reader can be defined that accurately interprets the text in the minimum number of saccades. Its computation uses three sources of information: 1) visual data, normally a few recognized letters in central vision and the locations of spaces in the periphery; 2) lexical data, including allowable words and their probabilities; and 3) eye-movement data, including distribution of saccade lengths.

Results from a computer simulation of the ideal reader may be informative about human readers. For example, the ideal reader exhibits regressive saccades (which also occur in human reading but are usually regarded as "errors") because ideal saccades of greatest expected length occasionally result in ambiguous interpretation of text. The ideal reader with scotomas has more regressions than normal and erratic eye movements (much larger standard deviation of saccade lengths), a pattern like that reported for some patients with central-field loss. The ideal reader is an explicit model for the combination of visual and other sources of information in reading. Its performance with abnormal retinal data may help us to understand the adverse effects of visual-field loss on human reading.

Psychophysics of Complex Auditory Signals. (Viemeister)

The major focus of the work during the past year has been on temporal aspects of auditory perception. This work was stimulated, in part, by our earlier work on "multiple looks" which provides an alternative to the notion of long time constant temporal

integration. The idea, essentially, is that we listen to the world through a brief (3-5 ms) temporal window and that we combine and selectively process information from these brief looks or samples. During the past year we performed an extensive detection theory analysis of auditory nerve recordings and showed that detection decisions based upon optimum combination of multiple looks, and decisions based upon true neural summation (integration), generally yield equivalent performance. Furthermore, the derived temporal integration functions were identical, demonstrating that a multiple look scheme can account for integration-like phenomena.

In another project, we extended the multiple look notion to the more realistic detection situation in which signals are presented at uncertain times. The experiment was a temporal analog of the probe-frequency method developed by Greenberg and Larkin: On 70% of the trials, so-called "primary" trials, the signal occurred in a fixed, well-marked temporal location; on the remaining trials the signal was presented in a random temporal location spanning a 1-sec. range. As expected, performance deteriorated as the probe location became more remote from that of the primary. Unexpectedly, the derived temporal window was quite broad, approximately 175 ms. We showed that this broad window was not the result of temporal uncertainty about the location of the primary. These results suggest certain weighting strategies that are used when signals are temporally uncertain. We are currently conducting an experimental investigation that is designed to further explore these possible strategies.

PUBLICATIONS AND SUBMITTED MANUSCRIPTS

Biederman, I., Gerhardstein, P. C., Cooper, E. E., & Nelson, C. A. (1992). High level object recognition without an inferior temporal lobe. Submitted to *Cognitive Neuroscience*.

Biederman, I., & Gerhardstein, P. C. (1992). Recognizing depth-rotated objects: Evidence for 3D viewpoint invariance. Submitted to *Journal of Experimental Psychology: Human Perception & Performance*.

Cooper, E. E., & Biederman, I. (1992). Priming objects with single volumes: Searching for the representational locus of perceptual priming. Submitted to *Acta Psychologica*.

Kersten, D., Bulthoff, H. H., Schwartz, B. & Kurtz, K. (1992). Interaction between transparency and structure from motion. *Neural Computation, 4*, 573-589.

Knill, D. C. (1992) The perception of surface contours and surface shape: from computation to psychophysics, *Journal of the Optical Society of Ameria, A* (in press).

Knill, D. C., Mamassian, P. & Kersten, D. (1992). The geometry of shadows. Submitted to *The International Journal of Computer Vision*.

Liu, Z., Kersten, D. & Knill, D. C. (1992). Object classification for human and ideal observers. Submitted to *Vision Research*.

Thompson, W. B., Kersten, D. & Knecht, W. R. (1992). Structure-from-motion based on information at surface boundaries. *Biological Cybernetics, 66*, 327-333.

Viemeister, N.F., Shivapuja, B.G., and Recio, A (1992). Physiological correlates of temporal integration. In Cazals, Y., Demany, L., and Horner, K. (Eds.) *Auditory Physiology and Perception*. Pergamon,
Oxford, pp. 323-329.

Wurm L.H., Legge, G.E., Isenberg, L.M. & Luebker A. (1992). Color improves object recognition in normal and low vision. Submitted to *J. Exp. Psychol.: Human Perception and Performance*.

CONFERENCE PRESENTATIONS AND ABSTRACTS

Chang, P. and Viemeister, N.F. (1991). Temporal windows for signals presented at uncertain times. *J. Acoust. Soc. Am., 90,* 2248.

Klitz, T.S., Mansfield, J.S., and Legge, G.E. (1992). Font popout in text images. *OSA Technical Digest*, 1992.

Legge, G.E. (1992). An ideal observer model of reading. ARVO 1992: Sarasota FL, *Suppl. Invest. Ophthalmol. & Vis. Sci., 33,* 3608.

Liu, Z., Kersten, D. & Knill, D. C. (1992). Object discrimination for human and ideal observers. ARVO 1992, Sarasota, Florida: The Association for Research in Vision and Ophthalmology, 825.

Madarasmi, S., Kersten, D. & Pong, T. C. (1992). A mult-layer approach to segmentation and interpolation with application to stereo vision. ARVO 1992, Sarasota, Florida: 1371.

Mamassian, P., Kersten, D. & Knill, D. C. (1992). Spatial layout from cast shadows. ARVO 1992, The Association for Research in Vision and Ophthalmology, 3191.

Mansfield, J.S., Akutsu, H., and Legge, G.E. (1992). Interocular contrast differences produce lateral shifts in the perceived location of binocular depth targets. ARVO 1992: *Suppl. Invest. Ophthalmol. & Vis. Sci., 33,* 1332.